Pranav Gangwar

| | |

Summary

Pursuing unanswerable questions at the intersection of computer engineering and biology, while judiciously leveraging today's computational power to push the boundaries of knowledge.

Education

2021 - Present PhD in Computer Engineering at **UC San Diego** (GPA: 4.0)
2021 - 2024 MS in Computer Engineering at **UC San Diego** (GPA: 4.0)

Research Projects

Wastewater based Epidemiology Turakhia Lab (Jan 2023 - Present)

- Developed a **scalable phylogenetic-based algorithm** that analyzes genomic sequences from wastewater samples to detect SARS-CoV-2 variants infecting the community.
- Reconstructs the genome sequences of all SARS-CoV-2 variants (among 16 million detected sequences) present in community wastewater, which contains very short genome fragments from multiple organisms.
- Provides significantly higher resolution by detecting SARS-CoV-2 variant sequences, compared to state-of-the-art software that can only identify the family of variants (among 3,000 variant families) infecting the community. This enhanced capability enables **early identification of emerging variants** and their outbreak clusters.
- Working with the **Centers for Disease Control and Prevention (CDC)** to analyze wastewater data across the United States with this algorithm and its potential integration to the National Wastewater Surveillance System (NWSS). Also working towards extending this algorithm to other diseases.

Machine Learning based pathogen evolution prediction Turakhia Lab (Jan 2024 - Present)

- Reformulated the SARS-CoV-2 evolution prediction problem by **integrating small stochastic variations** in natural selection and concentrating on key genomic regions. **Assessed learnability** with a Bayesian optimal classifier on a dataset of 16 million genome sequences.
- Comparing the performance of various machine learning models, along with integrating protein language models and evaluating their impact.

TermiNETor SEE LAB (Sep 2021 - Jun 2022)

- Designed a **bit-serial hardware architecture** to efficiently accelerate convolution operations in deep neural networks. Co-designed the architecture with a novel early convolution termination algorithm that conserves energy by predicting and eliminating redundant computations.
- Evaluated performance on **analytical simulator** and estimated area and power from RTL. Results showed 120x energy-efficiency over GPU and 2x over sparsity-aware accelerators.

Processing in memory based hardware accelerators SEE LAB (Sep 2021 - Dec 2022)

- Designed NeRF rendering low-power accelerator using **Processing-in-memory (PIM)** and **Near-data processing (NDP)**.
- Worked on DRAM based accelerator for fully homomorphic encryption based CKKS algorithm.

Publications

Minxuan Zhou, Yujin Nam, **Pranav Gangwar**, et al. (2023). *FHEmem: A Processing In-Memory Accelerator for Fully Homomorphic Encryption*. arXiv: 2311.16293 [cs.AR]. URL: https://arxiv.org/abs/2311.16293.

Uday Mallappa, **Pranav Gangwar**, Behnam Khaleghi, et al. (2022). "TermiNETor: Early Convolution Termination for Efficient Deep Neural Networks". In: *2022 IEEE 40th International Conference on Computer Design (ICCD)*, pp. 635–643. DOI: 10.1109/ICCD56317.2022.00098.

Pranav Gangwar, Satvik Maurya, Shubham Garg, et al. (2019). "Hardware/Software Co-Design of a High-Speed Othello Solver". In: *2019 IEEE 62nd International Midwest Symposium on Circuits and Systems (MWSCAS)*, pp. 1223–1226. DOI: 10.1109/MWSCAS.2019.8885136.

Pranav Gangwar, Neeta Pandey, and Rajeshwari Pandey (2019). "Novel Control Unit Design for a High-Speed SHA-3 Architecture". In: *2019 IEEE 62nd International Midwest Symposium on Circuits and Systems (MWSCAS)*, pp. 904–907. DOI: 10.1109/MWSCAS.2019.8885323.

Work Experience

GPU Architecture Intern - Intel Jun-Sept 2023, Jun-Sept 2022

- Modelled next generation Tensor cores on analytical simulator and explored different architectures
- Implemented dense matrix multiplication decomposition algorithm inspired from Stream-K, which picked best tile size based on Tensor core's utilization and memory bandwidth bottlenecks
- Compared Stream-K's performance results on Intel's next generation GPU with existing matrix decomposition algorithm

Digital Design Engineer - Texas Instruments July 2018 - July 2021

- Designed backend for wireless transceiver subsystems: floorplanning, module placement, clock tree synthesis, routing, and physical verification. Also helped with restructuring of pipeline in the modules to meet stringent input-output timing constraints
- Developed power recovery flow that reduced total design power by recommending standard cell resizing at sign-off stage. It took advantage of lower pessimism in timing constraints at sign-off as compared to place and route stage

Skills

Languages	C++, Verilog, Python, Tcl, Shell
Softwares	PyTorch, Innovus™, Tempus™, Voltus™, PVS, Assura®, Vivado®, MATLAB®, Virtuoso®

Mentoring

Summer 2024	Abhishikta Panja, Graduate Student	UC San Diego
Spring 2024 - Summer 2024	Manu Bhat, Undergraduate Student	UC San Diego
Spring 2024	Reya Sadhu, Graduate Student	UC San Diego
Spring 2023 - Winter 2024	Girish Krishna, Undergraduate Student	UC San Diego
Spring 2023 - Winter 2024	Carolyn Zhang, Undergraduate Student	UC San Diego